arXiv:1609.01100v1 [cs.CV] 5 Sep 2016

A max-cut approach to heterogeneity in cryo-electron microscopy

Yariv Aizenbud[1] Yoel Shkolnisky[1]

[1]School of Mathematical Sciences, Tel Aviv University, Israel

September 6, 2016

Abstract

The field of cryo-electron microscopy has made astounding advancements in the past few years, mainly due to improvements in the hardware of the microscopes. Yet, one of the key open challenges of the field remains the processing of heterogeneous data sets, produced from samples containing particles at several different conformational states. For such data sets, one must first classify their images into homogeneous groups, where each group corresponds to the same underlying structure, followed by reconstruction of a three-dimensional model from each of the homogeneous groups. This task has been proven to be extremely difficult. In this paper we present an iterative algorithm for processing heterogeneous data sets that combines the classification and reconstruction steps. We prove accuracy and stability bounds on the algorithm, and demonstrate it on simulated as well as experimental data sets.

Keywords: cryo-electron microscopy, single particle, three-dimensional reconstruction, heterogeneity, classification, max-cut, graph partitioning.

1 Introduction

The study of the molecular structure of complex proteins has drawn many efforts in the past few decades. Among the many structure determination methodologies available, cryo-electron microscopy (cryo-EM) single particle reconstruction (SPR) [7] has become a widespread and powerful tool [6], due to recent improvements in instrumentation as well as in the accompanying data processing algorithms [4]. These improvements resulted in three-dimensional molecular models with unprecedented resolutions as high as 2.2Å [5].

The process of resolving the three-dimensional structure of a molecule using cryo-EM SPR typically consists of the following steps. First, a sample consisting of many copies of the investigated

molecule is rapidly frozen and is imaged by an electron-microscope. This results in a large image of the sample, known as a micrograph, containing several dozens of particle images. The individual particle images are then segmented from this micrograph, resulting in a stack of images, where each image corresponds to a projection of one of the copies of the molecule in the sample. This process is repeated until a sufficiently large stack of raw images is obtained. The images in the stack are then clustered, aligned, and averaged, resulting in images of improved quality, known as class averages. The class averages are used for obtaining a low resolution model of the molecule, which is then refined using the stack of raw images into a high resolution model [8, 20].

Algorithms for estimating a low resolution model of the investigated molecule are often based on detecting common lines between pairs of images [22, 25, 27]. The underlying assumption of such algorithms is that all images were generated from exactly the same underlying molecule. Unfortunately, in many cases it is impossible to purify a sample consisting of only a single type molecule. In such cases another step is required, which classifies the images into groups such that all images in the same group correspond to the same underlying structure. This problem is known as the heterogeneity problem.

There were several previous attempts to address the heterogeneity problem. These attempts typically follow one of three approaches. The first approach consists of algorithms that are based on maximum likelihood estimation (MLE) [19]. The idea of such algorithms is to formulate a function that attains a minimum for the correct assignment of the images into the different classes, and to search for its minimum. As this function is highly nonlinear and very high dimensional, these algorithms critically depend on their initialization, converge only locally, and are overwhelmingly computationally expensive. The second approach is based on approximating the covariance between the voxels of the different volumes [13,14,17]. The resulting covariance matrix describes the differences between the different volumes. The drawback of this method is the high computational complexity needed to approximate the covariance matrix, which becomes prohibitive for volumes of high resolution. The third approach is based on building a similarity measure between pairs of images, and then dividing the images into different classes based on that similarity measure [12,21].

In this paper we propose an algorithm to address the heterogeneity problem that follows the third approach, and analyze its properties. In particular, we derive bounds on the accuracy of our algorithm. The presented algorithm takes N two-dimensional images generated from one of K different underlying structures, and returns a three-dimensional reconstruction for each of the K structures. We also show that the performance of the algorithm improves (under certain assumptions) not only as the quality of the images improves, but also as the number of images grows. The idea of the algorithm is to use a score to measure the similarity between every two images from the N input images. The important feature of the score is that it gives a low score to images of the same underlying structure, and gives a high score to images of different underlying structures. After calculating the score for every pair of images, we use a maximum k-cut algorithm to determine which of the two-dimensional images belong to each class. Unlike other works that used

this approach, in this paper we present an iterative algorithm that uses not only the information of common lines' correlations but also uses the estimated imaging directions and their consistency with the common lines.

The paper is organized as follows. In Section 2, we formulate the problem and set up the required notation. In Section 3, we review the required mathematical background. We describe the algorithm in details in Section 4, and prove its performance bounds in Section 5. We demonstrate the algorithm on simulated and experimental data in Section 6. Some concluding remarks are given in Section 7.

2 Problem formulation

We start by presenting the homogeneous setting of cryo-EM single particle reconstruction. Mathematically, a molecule is modeled as a function $\phi(x, y, z)$ describing the electric potential of the molecule. Under an ideal model, each image P_i generated by the electron microscope is given by

$$P_i(x, y) = \int_{-\infty}^{\infty} \phi\left(R_i r\right)\, dz, \quad r = (x, y, z)^T, \quad i = 1, \ldots, N, \tag{1}$$

where $R_i \in SO(3)$ is an unknown rotation, which describes the orientation of the molecule ϕ at the moment of imaging. In this (homogeneous) setting, all images are assumed to correspond to exactly the same underlying molecule ϕ. The goal of cryo-EM structure determination is to recover ϕ given a finite set of its images $P_1, \ldots, P_N$ generated according to (1).

In the heterogeneous setting of cryo-EM structure determination, we have K underlying molecules $\phi_1, \ldots, \phi_K$, and each 2D image P_i is generated according to (1) from one of $\phi_1, \ldots, \phi_K$. Specifically, if image P_i was generated from ϕ_k according to (1), then we denote $C(P_i) = k$, that is, the "class" of image P_i is k. We denote by G_k the indices of all images that correspond to the molecule ϕ_k, namely

$$G_k = \left\{i \,|\, C(P_i) = k\right\}, \quad k = 1, \ldots, K. \tag{2}$$

Thus, $G_1, \ldots, G_K$ are disjoint sets whose union is equal to the set $\{1, \ldots, N\}$. Our goal is to estimate the sets G_k, $k = 1, \ldots, K$, and the rotations R_i, $i = 1, \ldots, N$. Once these sets and rotations are estimated, the structures $\phi_1, \ldots, \phi_K$ can be reconstructed using standard algorithms [11, 16]. In this work we assume that K is known.

3 Partitioning a graph: max-cut

In this section we present the max-cut problem, its Goemans-Williamson approximation algorithm, and a particular case for which this algorithm is exact. This particular case is used later for the analysis of our algorithm.

Let (V, E) be a weighted graph, with a set of vertices V and a set of edges E. We denote the nodes in V by $v_1, \ldots, v_N$, so that $|V| = N$. We assume that each edge $(v_i, v_j) \in E$ is associated with a real-valued positive weight $w(v_i, v_j)$. Whenever $(v_i, v_j) \notin E$, we set $w(v_i, v_j) = 0$. The adjacency matrix W of the graph (V, E) is a matrix of size $N \times N$ with entries $w_{ij} = w(v_i, v_j)$. For simplicity of notation, we identify the set $\{v_1, \ldots, v_N\}$ with the set $\{1, \ldots, N\}$ (via the trivial map $v_i \mapsto i$). A cut (sometimes called a 2-cut) is a partition of the set of vertices V into two disjoint subsets , namely, into G_1 and G_2 such that $G_1 \cup G_2 = \{1, \ldots, N\}$ and $G_1 \cap G_2 = \emptyset$. The weight of a cut is defined as

$$W(G_1, G_2) = \sum_{i \in G_1, j \in G_2} w_{ij}. \tag{3}$$

The maximum-cut (max-cut) problem is to find a cut whose weight is maximal among all possible cuts. Although this problem has been proven to be NP-complete, it has many approximation algorithms. The currently best approximation algorithm (and assuming the unique games conjecture, also the best approximation possible with polynomial complexity) is a randomized algorithm by Goemans and Williamson [10]. This algorithm guarantees that the value of its returned cut is at least 0.87 of the optimal result.

There is an analogous formulation to the max-cut problem for partitioning a graph into K sub-graphs, which is called max K-cut. In this case, a K-cut is a partition of V into K disjoint sets, that is, into the sets $G_1, \ldots, G_K$ such that

$$\bigcup_{k=1}^{K} G_k = V, \quad G_i \cap G_j = \emptyset, \ i \neq j.$$

The weight of the cut in this case is defined as

$$W(G_1, \ldots, G_K) = \sum_{k=1}^{K} \sum_{i \in G_k, j \notin G_k} w_{ij}, \tag{4}$$

that is, the weight of the cut is the sum of all the edges whose endpoint vertices are in two different sets. Similar to the Goemans-Williamson algorithm, there are algorithms [9] that can be applied in this case, which result, for some values of K, in even better approximation bounds than for the $K = 2$ case.

Although the Goemans-Williamson algorithm finds only an approximate solution, there is a special family of graphs for which its solution is exact.

Definition 1. *The graph (V, E) is bipartite if its vertices can be divided into two disjoint sets such that every edge connects a vertex in the first set to one in the second set.*

We next show in Lemma 2 that for a bipartite graph, the Goemans-Williamson algorithm finds the optimal cut (and not an approximation of it). This lemma is used later to analyze the

properties of our algorithm.

Lemma 2. *The Goemans-Williamson algorithm finds the exact solution for the max-cut problem of a bipartite graph.*

Proof. As shown in [10], the solution to the max-cut problem is given by the rank-1 matrix Σ that minimizes trace$(W\Sigma)$ such that Σ is positive semidefinite, $\Sigma = \sigma\sigma^T$ where $\sigma = (\sigma_1, \ldots, \sigma_N)^T$, and $\sigma_i \in \{+1, -1\}$. The sign of σ_i represents the subset of the cut in which vertex i is included. The Goemans-Williamson algorithm discards the rank-1 assumption, that is, minimizes trace$(W\Sigma)$ for Σ positive semidefinite with $\sigma_{ii} = 1$.

Without loss of generality, assume that the adjacency matrix of the bipartite graph (V, E), denoted by W, is a block matrix with the structure

$$W = \begin{pmatrix} 0_{N_1 \times N_1} & A_{N_1 \times N_2} \\ A^T_{N_2 \times N_1} & 0_{N_2 \times N_2} \end{pmatrix}. \tag{5}$$

Since $\sigma_{ij} \in \{+1, -1\}$, if we denote the elements of A in (5) by a_{ij}, we have that

$$a = -2 \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} a_{ij} \leq \text{trace}(W\Sigma). \tag{6}$$

Any cut G_1, G_2 of a graph (V, E) can be encoded as a vector consisting of $+1$ and -1, whose i'th coordinate equals 1 if node i is in G_1, and equals -1 if node i is in G_2. In the case of the matrix W in (5), we define σ^{opt} to consist of a block of 1's of length N_1 followed by a block of -1's of length N_2. It can be easily verified that the cut encoded by σ^{opt} achieves the bound a in (6) and is thus optimal.

We denote $\bar{\Sigma} = \sigma^{\text{opt}}\sigma^{\text{opt}^T}$. We now show that $\bar{\Sigma}$ is the minimum of the Goemans-Williamson optimization problem (i.e., without the rank-1 restriction). Since the Goemans-Williamson optimization problem optimizes over positive semidefinite matrices Σ, all the 2×2 main minors of Σ are non-negative, that is, $\det \Sigma_{ij} \geq 0$, where

$$\Sigma_{ij} = \begin{pmatrix} \sigma_{ii} & \sigma_{ij} \\ \sigma_{ji} & \sigma_{jj} \end{pmatrix}.$$

In other words, $\sigma_{ii}\sigma_{jj} - \sigma_{ij}\sigma_{ji} \geq 0$. Since $\sigma_{ii} = 1$ we get that $1 - \sigma_{ij}\sigma_{ji} \geq 0$ or, using the symmetry of Σ, $1 - \sigma_{ij}^2 \geq 0$. Thus, $|\sigma_{ij}| \leq 1$. For matrices Σ of size $(N_1 + N_2) \times (N_1 + N_2)$ we have that

$$\text{trace}(W\Sigma) = \sum_{i=1}^{N_1+N_2} \sum_{j=1}^{N_1+N_2} w_{ij}\sigma_{ji} = \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} a_{ij}\sigma_{N_1+j,i} + \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} a_{ji}\sigma_{j,N_1+i}. \tag{7}$$

Since the entries of Σ are less than or equal to 1, the least possible value of (7) is $-2\sum_{i=1}^{N_1}\sum_{j=1}^{N_2}(a_{ij})$

which is achieved by $\bar{\Sigma}$. $\square$

4 Algorithm description

Let P_i and P_j be two images generated according to (1) (using the same ϕ). If we denote by $\hat{P}_i$ the two-dimensional Fourier transform of P_i and by $\hat{\phi}$ the three-dimensional Fourier transform of ϕ, then, the Fourier projection-slice theorem [15] implies that $\hat{P}_i$ is the restriction of $\hat{\phi}$ to the plane spanned by the first two columns of R_i of (1). Explicitly,

$$\hat{P}_i(\omega_x, \omega_y) = \hat{\phi}\left(\omega_x R_i^{(1)} + \omega_y R_i^{(2)}\right), \tag{8}$$

where $R_i^{(1)}$, $R_i^{(2)}$, $R_i^{(3)}$ are the columns of the rotation matrix R_i. As a consequence of (8), any two (Fourier-transformed) images $\hat{P}_i$ and $\hat{P}_j$ share a common line through the origin, namely, there exist direction vectors $c_{ij}, c_{ji} \in \mathbb{R}^2$ such that $\hat{P}_i(\xi c_{ij}) = \hat{P}_j(\xi c_{ji})$ for any $\xi \in \mathbb{R}$. The vectors c_{ij} and c_{ji} are given explicitly by [23]

$$c_{ij} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} R_i^T \frac{R_i^{(3)} \times R_j^{(3)}}{\|R_i^{(3)} \times R_j^{(3)}\|}, \quad c_{ji} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix} R_j^T \frac{R_i^{(3)} \times R_j^{(3)}}{\|R_i^{(3)} \times R_j^{(3)}\|}. \tag{9}$$

If we further lift the vectors c_{ij} and c_{ji} to $\mathbb{R}^3$ by zero padding, then it can be shown that for all i and j it holds that $R_{ij}c_{ij} = R_{ji}c_{ji}$ (see [22] for a detailed proof). Now assume that we are given rotations $\tilde{R}_i$ and $\tilde{R}_j$, which are estimates of R_i and R_j, as well as vectors $\tilde{c}_{ij}$ and $\tilde{c}_{ji}$, which are estimates for c_{ij} and c_{ji}. Then, we define a similarity score for any pair of images i and j by $\|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\|$. This score is 0 if the common lines and the rotations are correct, and is small for small errors in the common lines or in the rotations. The LUD algorithm [27] finds rotations $\tilde{R}_i$, $i = 1, \ldots, N$ that bring the score $\sum_{i,j} \|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\|$ to a local minimum.

As the algorithm [27] assumes the homogeneous setting, namely, that all images were generated from the same underlying ϕ, it cannot be applied directly to the heterogeneous setting. However, it can be applied to each of the (unknown) sets G_k of (2).

Consider a graph (V, E) whose i'th vertex corresponds to image P_i, and where any two vertices are connected by an edge. Our goal is to partition the vertices of the graph (V, E) into the sets G_k of (2). Assume moreover that we are given weights for the edges E, and denote by w_{ij} the weight of the edge between vertices i and j. In this case, according to (4), the weight of some cut

$\tilde{G} = \left\{\tilde{G}_1, \ldots, \tilde{G}_K\right\}$ is equal to

$$\begin{aligned} W(\tilde{G}) &= \sum_{k=1}^{K} \sum_{i\in\tilde{G}_k, j\notin\tilde{G}_k} w_{ij} = \sum_{k=1}^{K} \left(\sum_{\substack{j=1 \\ i\in\tilde{G}_k}}^{N} w_{ij} - \sum_{i,j\in\tilde{G}_k} w_{ij} \right) \\ &= \sum_{i,j=1}^{N} w_{ij} - \sum_{k=1}^{K} \sum_{i,j\in\tilde{G}_k} w_{ij} = C - \sum_{k=1}^{K} \sum_{i,j\in\tilde{G}_k} w_{ij}, \end{aligned} \tag{10}$$

where C is the sum of all weights in the graph (independent of the partition). Thus, finding the cut that maximizes $W(\tilde{G})$ is equivalent to finding the cut minimizing $\sum_{k=1}^{K} \sum_{i,j\in\tilde{G}_k} w_{ij}$. If we now set

$$w_{ij} = \|R_i c_{ij} - R_j c_{ji}\|, \tag{11}$$

we get that $W(\tilde{G}) \geq 0$ for any partition $\tilde{G}$, and that $W(G_1, \ldots, G_K) = 0$. In other words, the rotations $R_1, \ldots, R_N$ and the partition $G = \{G_1, \ldots, G_K\}$ are obtained as the solution to the optimization problem

$$\min_{\tilde{G},\tilde{R}} \sum_{k=1}^{K} \sum_{i,j\in\tilde{G}_k} \|\tilde{R}_i c_{ij} - \tilde{R}_j c_{ji}\|. \tag{12}$$

Since the optimization problem in (12) is high-dimensional and non-convex, we propose the following descend procedure to find its minimum. Start from an initial estimate for R and G, and minimize (12) by minimizing alternatingly over $\tilde{R}$ and $\tilde{G}$ as follows:

1. Find the minimizing $\tilde{R}$ for a given partition $\tilde{G}$ (known from the previous iteration), by applying the LUD algorithm [27] on each of $\tilde{G}_k$, $k = 1, \ldots, K$.

2. Given the new rotations, find a partition $\tilde{G}_1, \ldots, \tilde{G}_K$. This can be achieved approximately using the Goemans-Williamsons algorithm [10].

Once the rotations $\tilde{R}$ have been estimated in step 1, minimizing (12) in step 2 is equivalent to maximizing (10) with weights given by (11). However, the objective in (10) is exactly the one maximized by the solution to the max-cut problem applied to the graph (V, E) defined above, with weights given by (11). Optimization algorithms for steps 1 and 2 (minimizations over the rotations and the partitions) provide only approximate solutions, and thus, the solution to (12) is also only approximate. Nevertheless, we show later (Theorem 8) that in some cases we can bound the quality of this approximate solution. A detailed description of the algorithm is given in Algorithm 1.

Algorithm 1 Reconstruction from heterogeneous data sets

Input: c_{ij} Common lines between all images $P_1, \ldots, P_N$.
K Number of groups.
Output: $\tilde{R}_1, \ldots, \tilde{R}_N$ Estimated rotations of all the images.
$\tilde{G}_1, \ldots, \tilde{G}_K$ Estimated partition of the images into homogeneous groups.
▷ Initialization
$\tilde{G}_1^{(1)} \leftarrow \{1, \ldots, N\}$ ▷ Start from an initial guess. For example, all images in $\tilde{G}_1^{(1)}$.
for $k = 2$ to K **do**
  $\tilde{G}_k^{(1)} \leftarrow \emptyset$
end for
$n \leftarrow 0$ ▷ Iteration number.
repeat
  $n \leftarrow n + 1$
  ▷ LUD step
  for $k = 1$ to K **do**
    $C_k^{(n)} = \left\{ c_{ij} \mid i \in \tilde{G}_k^{(n)} \wedge j \in \tilde{G}_k^{(n)} \right\}$ ▷ Construct common lines matrix for $\tilde{G}_k^{(n)}$.
    $\{\tilde{R}_i^{(n)} | i \in \tilde{G}_k^{(n)}\} = \mathrm{LUD}(C_k^{(n)})$ ▷ Find rotations for images in $\tilde{G}_k^{(n)}$ using LUD [27].
    $\tilde{R}^{(n)} = \{\tilde{R}_1^{(n)}, \ldots, \tilde{R}^{(n)}\}$
    if $F(\tilde{R}^{(n)}, \tilde{G}^{(n-1)}) > F(\tilde{R}^{(n-1)}, \tilde{G}^{(n-1)})$ **then** ▷ Where F is defined in (13).
      $\tilde{R}^{(n)} = \tilde{R}^{(n-1)}$
    end if
  end for
  ▷ max-cut step
  for $i = 1$ to N **do**
    for $j = 1$ to N **do**
      $W_{ij} \leftarrow \|\tilde{R}_i c_{ij} - \tilde{R}_j c_{ji}\|$ ▷ W is the graph to partition.
    end for
  end for
  $\left[\tilde{G}_1^{(n)}, \ldots, \tilde{G}_K^{(n)}\right] \leftarrow \text{max-cut}(W, K)$ ▷ Apply max K-cut.
  $\tilde{G}^{(n)} = \{\tilde{G}_1^{(n)}, \ldots, \tilde{G}^{(n)}\}$
  if $F(\tilde{R}^{(n)}, \tilde{G}^{(n)}) > F(\tilde{R}^{(n)}, \tilde{G}^{(n-1)})$ **then**
    $\tilde{G}^{(n)} = \tilde{G}^{(n-1)}$
  end if
until $F(\tilde{R}^{(n)}, \tilde{G}^{(n)}) = F(\tilde{R}^{(n-1)}, \tilde{G}^{(n-1)})$ ▷ The algorithm converges.
return $\tilde{R}_1^{(n)}, \ldots, \tilde{R}_N^{(n)}$ and $\tilde{G}_1^{(n)}, \ldots, \tilde{G}_K^{(n)}$

5 Convergence and error bounds

For simplicity of the presentation, we assume in this section that $K = 2$, as all arguments are easily extended to any K. We will show that Algorithm 1 converges, and that it is stable around its optimum, namely, that when the initial partition and rotations are close to the optimum, step 23 of Algorithm 1 finds a good class partition with high probability.

We start by establishing the convergence of Algorithm 1.

Theorem 3. *Algorithm 1 converges.*

Proof. We denote the objective function from (12) by

$$F(R, G) = \sum_{k=1}^{2} \sum_{i,j \in G_k} \|R_i c_{ij} - R_j c_{ji}\|. \tag{13}$$

Since in each iteration of Algorithm 1 we do not increase the value of $F(R, G)$, we have that $F(R^{(n)}, G^{(n)}) \geq F(R^{(n+1)}, G^{(n+1)})$, and thus $F(R^{(n)}, G^{(n)})$ is monotonically non-increasing and bounded by zero. □

Next, we derive error bounds for Algorithm 1. We start by proving that if the common lines are detected correctly, and the rotations are assigned correctly, then the algorithm finds the correct partition $\{G_1, G_2\}$.

Theorem 4. *Suppose that the common lines $c_{ij} \in \mathbb{R}^2$ in Algorithm 1 are correct when the images P_i and P_j are in the same class, and are uniformly distributed otherwise. If the rotations $R_1, \ldots, R_N$ in step 12 of Algorithm 1 are assigned correctly, then Algorithm 1 will find the correct class partitioning when executing step 23.*

Proof. For the correct rotations $R_1, \ldots, R_N$, we have that for any two images P_i and P_j in the same class it holds that $\|R_i c_{ij} - R_j c_{ji}\| = 0$. For P_i and P_j in different classes, the score $\|R_i c_{ij} - R_j c_{ji}\|$ is a random variable that gets the value 0 with probability 0. Thus, if we consider the graph whose adjacency matrix W is given by $W_{ij} = \|R_i c_{ij} - R_j c_{ji}\|$, we get a bipartite graph, since all the images from the same class are not connected and images from different classes are connected with some random weight. By Lemma 2, the Goemans-Williamson algorithm returns the correct partition for this graph. □

In order to show that for common lines and rotations "close" to the correct ones, Algorithm 1 still finds a partition that is "close" to the correct one, we need to define what are "close" common lines and rotations, as well as what are "close" partitions.

Definition 5. *Let c_{ij} and c_{ji} be the correct common line between images P_i and P_j, as defined by (9). Let $\tilde{c}_{ij}$ and $\tilde{c}_{ji}$ be some estimate of c_{ij} and c_{ji}, respectively. We define the distance between*

c_{ij} and $\tilde{c}_{ij}$ as the angle between them, namely,

$$d(c_{ij}, \tilde{c}_{ij}) = \arccos\langle c_{ij}, \tilde{c}_{ij}\rangle, \tag{14}$$

where $\langle \cdot, \cdot \rangle$ is the standard dot product in $\mathbb{R}^2$.

Unless otherwise stated, all norms below refer to the induced 2-norm $\|\cdot\|_2$.

Lemma 6. *If $\|\tilde{R}_1 - R_1\| < \varepsilon$ and $\|\tilde{R}_2 - R_2\| < \varepsilon$, than $\|\tilde{R}_1\tilde{R}_2 - R_1R_2\| < 2\varepsilon$.*

Proof. From the triangle inequality and the fact that $\|R\| = 1$ for any $R \in SO(3)$, we get that $\|\tilde{R}_1\tilde{R}_2 - R_1R_2\| = \|\tilde{R}_1\tilde{R}_2 - \tilde{R}_1R_2 + \tilde{R}_1R_2 - R_1R_2\| \leq \|\tilde{R}_1(\tilde{R}_2 - R_2)\| + \|(\tilde{R}_1 - R_1)R_2\| < 2\varepsilon$. □

Next, we define the quality of a partition of a heterogeneous data set.

Definition 7. *A partition of $P_1, \ldots, P_N$ into the sets $\tilde{G}_1$, $\tilde{G}_2$ is called p-precise for class k, $k = 1, 2$, if $\left|\tilde{G}_k \cap G_k\right| / \left|\tilde{G}_k\right| \geq p$, where G_k is defined in (2), and $|G|$ is the number of elements in the set G. A partition is p-precise if it is p-precise for all k.*

In other words, a partition of the images $P_1, \ldots, P_N$ is p-precise if the ratio between the number of correct images in class k and the total number of images assigned to class k is at least p. Note that the correct class partition is 1-precise, and a random class partition is about $\frac{1}{K}$-precise.

Theorem 8 below shows that step 23 of Algorithm 1 gives "good" results if we start from a "good" initial state. For ease of notation, Theorem 8 is proven for the case $K = 2$, and we moreover assume that the correct partition satisfies $|G_1| = |G_2| = N/2$.

Theorem 8. *Let $P_1, \ldots, P_N$ be N images comprising a heterogeneous data set corresponding to $K = 2$. Let c_{ij} be the common line between P_i and P_j, and let $\tilde{c}_{ij}$ be some estimate of c_{ij} used as the input to Algorithm 1. Also, let R_i be the rotation corresponding to P_i (see (1)), and let $\tilde{R}_i$ be the rotation estimated in step 12 of Algorithm 1. Let $\varepsilon > 0$ and assume that*

1. *$\|R_i - \tilde{R}_i\| \leq \varepsilon$, $i = 1, \ldots, N$.*
2. *If $C(P_i) = C(P_j)$ then $d(c_{ij}, \tilde{c}_{ij}) \leq \varepsilon$, $i, j = 1, \ldots, N$.*
3. *If $C(P_i) \neq C(P_j)$ then $\tilde{c}_{ij}$ is uniformly distributed and independent of the rotations, $i, j = 1, \ldots, N$.*

Then, for sufficiently large N, with high probability, step 23 in Algorithm 1 results in at least $0.87 - \frac{63}{4}\varepsilon$-precise partition (according to Definition 7).

Assumption 1 in Theorem 8 states that the algorithm finds rotations that are close to the true ones. Assumption 2 states that common lines between images of the same class are detected with small error. Assumption 3 states that common lines between images of different classes (that therefore have no common line) are random, namely, we don't make consistent errors when detecting common lines that do not exist.

Proof. Let (V, E) be a graph whose vertex $v_i \in V$ corresponds to image P_i, $i = 1, \ldots, N$, and whose (undirected) edge $(v_i, v_j) \in E$ has weight $\|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\|$. If P_i and P_j are in the same class, namely, $C(P_i) = C(P_j)$, we denote $\epsilon_{ij} = \|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\|$. Then,

$$\begin{aligned}
\|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\| &= \|\tilde{c}_{ij} - \tilde{R}_i^{-1}\tilde{R}_j\tilde{c}_{ji}\| \\
&\leq \|\tilde{c}_{ij} - c_{ij}\| + \|c_{ij} - \tilde{R}_i^{-1}\tilde{R}_j\tilde{c}_{ji} + \tilde{R}_i^{-1}\tilde{R}_j c_{ji} - \tilde{R}_i^{-1}\tilde{R}_j c_{ji}\| \\
&\leq \|\tilde{c}_{ij} - c_{ij}\| + \|c_{ij} - (\tilde{R}_i^{-1}\tilde{R}_j)(R_i^{-1}R_j)^{-1}(R_i^{-1}R_j)c_{ji}\| + \|\tilde{R}_i^{-1}\tilde{R}_j(\tilde{c}_{ji} - c_{ji})\| \quad (15) \\
&= \|\tilde{c}_{ij} - c_{ij}\| + \|c_{ij} - (\tilde{R}_i^{-1}\tilde{R}_j)(R_i^{-1}R_j)^{-1}c_{ij}\| + \|\tilde{c}_{ji} - c_{ji}\| \\
&\leq 4\varepsilon,
\end{aligned}$$

where the last inequality follows from assumptions 1 and 2 in Theorem 8, together with Lemma 6. In other words, $0 \leq \epsilon_{ij} \leq 4\varepsilon$.

Next, if $C(P_i) \neq C(P_j)$, we denote $X_{ij} = \|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\|$. Also, we denote $X'_{ij} = \|R_i\tilde{c}_{ij} - R_j\tilde{c}_{ji}\|$, that is, X'_{ij} is defined using the correct rotations. Note that both X_{ij} and X'_{ij} are random variables. Intuitively, we are going to show that X'_{ij} is "much larger" than ϵ_{ij} with high probability, and that X_{ij} is "close" to X'_{ij}. Thus, to maximize the cut of the graph (V, E), images P_i and P_j of the same class (satisfying $C(P_i) = C(P_j)$) for which ϵ_{ij} is small, should be assigned to the same subset of the partition.

Since we assume that the common lines $\tilde{c}_{ij}$ for i and j such that $C(P_i) \neq C(P_j)$ are uniformly random and independent of the rotations, the weights X'_{ij} are i.i.d and distributed as the distance between two random unit vectors in $\mathbb{R}^3$, whose distribution is specified in Lemma 9 in Appendix A.

Note that,

$$X_{ij} = \|\tilde{R}_i\tilde{c}_{ij} - \tilde{R}_j\tilde{c}_{ji}\| = \|\tilde{R}_i\tilde{c}_{ij} - R_i\tilde{c}_{ij} + R_i\tilde{c}_{ij} - R_j\tilde{c}_{ji} + R_j\tilde{c}_{ji} - \tilde{R}_j\tilde{c}_{ji}\|.$$

and

$$\begin{aligned}
\|R_i\tilde{c_{ij}} - R_j\tilde{c}_{ji}\| - \|\tilde{R}_i\tilde{c}_{ij} - R_i\tilde{c}_{ij}\| - \|R_j\tilde{c}_{ji} - \tilde{R}_j\tilde{c}_{ji}\| \leq \\
\leq \|\tilde{R}_i\tilde{c}_{ij} - R_i\tilde{c}_{ij} + R_i\tilde{c}_{ij} - R_j\tilde{c}_{ji} + R_j\tilde{c}_{ji} - \tilde{R}_j\tilde{c}_{ji}\| \leq \\
\leq \|R_i\tilde{c}_{ij} - R_j\tilde{c}_{ji}\| + \|\tilde{R}_i\tilde{c}_{ij} - R_i\tilde{c}_{ij}\| + \|R_j\tilde{c}_{ji} - \tilde{R}_j\tilde{c}_{ji}\|.
\end{aligned}$$

From assumption 1 in Theorem 8 it follows that $\|\tilde{R}_i\tilde{c}_{ij} - R_i\tilde{c}_{ij}\| = \|(\tilde{R}_i - R_i)\tilde{c}_{ij}\| \leq \|\tilde{R}_i - R_i\| \leq \varepsilon$, and thus we have

$$X'_{ij} - 2\varepsilon \leq X_{ij} \leq X'_{ij} + 2\varepsilon. \quad (16)$$

Next, we analyze the score of an arbitrary partition of V and show that with high probability the score of "bad" partitions is low and of "good" partitions is high. Thus, we get a bound for how "bad" the partition generated by Algorithm 1 can be. Let G_1 and G_2 be the sets defined in (2),

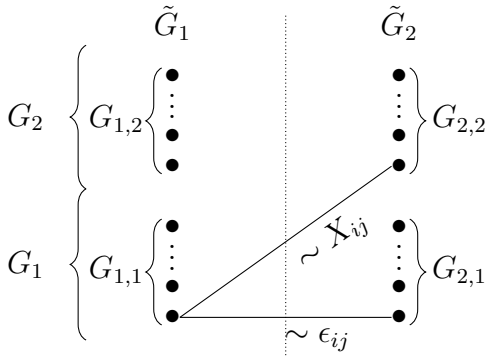

Figure 1: Calculating the weight of a cut $\tilde{G}_1$, $\tilde{G}_2$. The weight of an edge between vertex i in $G_{1,1}$ and vertex j in $G_{2,1}$ is the random variable ϵ_{ij}. The weight of an edge between vertex i in $G_{1,1}$ and vertex j in $G_{2,2}$ is the random variable X_{ij}.

and let $\tilde{G}_1$ and $\tilde{G}_2$ be some partition of the graph (V, E). Also, recall that we assume that $K = 2$ and that $|G_1| = |G_2| = N/2$. We denote $M = N/2$. Intuitively, G_1 and G_2 is the ground truth partition of the graph, and $\tilde{G}_1$ and $\tilde{G}_2$ is the cut returned by our algorithm. We denote

$$G_{k,l} = \tilde{G}_k \cap G_l, \quad k, l = 1, 2, \tag{17}$$

and so

$$\begin{aligned} |G_{1,1}| &= \left|\tilde{G}_1 \cap G_1\right| = p_1 M, & |G_{1,2}| &= \left|\tilde{G}_1 \cap G_2\right| = p_2 M, \\ |G_{2,1}| &= \left|\tilde{G}_2 \cap G_1\right| = (1-p_1) M, & |G_{2,2}| &= \left|\tilde{G}_2 \cap G_2\right| = (1-p_2) M, \end{aligned} \tag{18}$$

where $0 \le p_1, p_2 \le 1$ and we have used the fact that $G_1 \cup G_2 = V$ and $G_1 \cap G_2 = \emptyset$. We thus get that

$$\left|\tilde{G}_1\right| = (p_1 + p_2) M, \quad \left|\tilde{G}_2\right| = (2 - p_1 - p_2) M.$$

Using the notation of (17), $\tilde{G}_1$ is our "estimate" for G_1, and the subsets $G_{1,1}$ and $G_{1,2}$ are the subsets of $\tilde{G}_1$ that were assigned "correctly" and "incorrectly", respectively. The case for $G_{2,1}$ and $G_{2,2}$ is analogous.

The weight of a cut $\tilde{G}_1$ and $\tilde{G}_2$ corresponding to parameters p_1 and p_2, respectively, can be expressed in terms of X_{ij} and ϵ_{ij} as follows. It equals to the sum of weights of edges between $G_{1,1}$ and $G_{2,1}$ (with weight ϵ_{ij}), between $G_{1,1}$ and $G_{2,2}$ (with weight X_{ij}), between $G_{1,2}$ and $G_{2,1}$ (with weight X_{ij}), and between $G_{1,2}$ and $G_{2,2}$ (with weight ϵ_{ij}). A graphical illustration of this setup is

given in Figure 1. Formally, by denoting $S(p_1, p_2) = W(\tilde{G}_1, \tilde{G}_2)$ (see (3)), we have that

$$\begin{aligned} S(p_1,p_2) &= \sum_{i\in G_{1,1}} \left(\sum_{j\in G_{2,1}} \epsilon_{ij} + \sum_{j\in G_{2,2}} X_{ij} \right) + \sum_{i\in G_{1,2}} \left(\sum_{j\in G_{2,2}} \epsilon_{ij} + \sum_{j\in G_{2,1}} X_{ij} \right) \\ &= \sum_{(i,j)\in B_X} X_{ij} + \sum_{(i,j)\in B_\epsilon} \epsilon_{ij} \qquad (19) \\ &\leq \sum_{(i,j)\in B_X} (X'_{ij} + 2\varepsilon) + |B_\epsilon| 4\varepsilon, \qquad (20) \end{aligned}$$

where

$$B_X = (G_{1,1} \times G_{2,2}) \cup (G_{1,2} \times G_{2,1}), \quad B_\epsilon = (G_{1,1} \times G_{2,1}) \cup (G_{1,2} \times G_{2,2}), \qquad (21)$$

and (20) was derived using (15) and (16). Note that

$$|B_X| = M^2(p_1 + p_2 - 2p_1p_2), \quad |B_\epsilon| = M^2(p_1(1-p_1) + p_2(1-p_2)). \qquad (22)$$

By (19) and (16), the value of $S(p_1, p_2)$ for $p_1 = 0$ and $p_2 = 1$ (or symmetrically for $p_1 = 1$ and $p_2 = 0$) satisfies

$$S(0,1) \geq \left[\sum_{(i,j)\in B_X} X'_{ij} \right] - 2\varepsilon M^2, \qquad (23)$$

where we have used the fact that in this case $|B_X| = M^2$ and $|B_\epsilon| = 0$. The expected value of the right hand side of (23) is $M^2E(X') - 2\varepsilon M^2$, where X' denotes any of the i.i.d random variables X'_{ij}, and its variance is $M^2Var(X')$. By Chebyshev's inequality,

$$P\left(S(0,1) < M^2E(X') - 2\varepsilon M^2 - kM\sigma(X')\right) < \frac{1}{k^2}. \qquad (24)$$

This means that for large enough k, with high probability,

$$S(0,1) \geq M^2E(X') - 2\varepsilon M^2 - kM\sigma(X'), \qquad (25)$$

that is, we have a lower bound for the weight of the correct partition.

Next, we denote by $U \subset [0,1]^2$ the set of (p_1, p_2) such that p_1M and p_2M are integers, and for any $0 \leq \delta \leq 1$, we denote by U_δ the subset of U such that $p_1 \geq \delta, 1 - p_2 \geq \delta$ or $1 - p_1 \geq \delta, p_2 \geq \delta$ (each $(p_1, p_2) \in U_\delta$ corresponds to a partition that is δ-precise in (18)).

In order to ensure that step 23 of Algorithm 1 returns a δ-precise partition with high probability, it is enough to require that with high probability the maximal score of a non-δ-precise partition is less than 0.87 of the correct partition, that is,

$$P\left(0.87S(0,1) < \max_{(p_1,p_2)\in U\setminus U_\delta} S(p_1,p_2)\right) \ll 1, \qquad (26)$$

or, alternatively, we will find the maximal δ such that with high probability

$$0.87S(0,1) \geq \max_{(p_1,p_2)\in U\backslash U_\delta} S(p_1,p_2). \tag{27}$$

We next estimate the maximum weight over all the cuts with $(p_1,p_2) \in U\backslash U_\delta$. Since there are 2^{2M} possible cuts $(\tilde{G}_1,\tilde{G}_2)$, we have that $|U\backslash U_\delta| \leq |U| = 2^{2M}$ (note that we assumed that $|G_1| = |G_2| = M$, but Algorithm 1 can return $\tilde{G}_1$ and $\tilde{G}_2$ of any size). Using (20) and Lemma 10 in Appendix A we get that with high probability

$$\begin{aligned}
\max_{(p_1,p_2)\in U\backslash U_\delta} S(p_1,p_2) &< \max_{(p_1,p_2)\in U\backslash U_\delta} \sum_{(i,j)\in B_X} (X'_{ij}+2\varepsilon) + |B_\epsilon|4\varepsilon \\
&< \max_{(p_1,p_2)\in U\backslash U_\delta} M^2\left[(p_1+p_2-2p_1p_2)(E(X')+2\varepsilon)+(p_1(1-p_1)+p_2(1-p_2))4\varepsilon\right] \\
&\quad +2\log(|U\backslash U_\delta|)\left[(p_1+p_2-2p_1p_2)\sigma(X')\right] \\
&< \max_{(p_1,p_2)\in U\backslash U_\delta} M^2\left[(p_1+p_2-2p_1p_2)(E(X')+2\varepsilon)+(p_1(1-p_1)+p_2(1-p_2))4\varepsilon\right] \\
&\quad +2\sqrt{2M}\left[(p_1+p_2-2p_1p_2)\sigma(X')\right].
\end{aligned} \tag{28}$$

If we now require that

$$\begin{aligned}
&0.87\left(M^2E(X')-2\varepsilon M^2-kM\sigma(X')\right) \geq \\
&\max_{(p_1,p_2)\in U\backslash U_\delta} M^2[(p_1+p_2-2p_1p_2)(E(X')+2\varepsilon)+(p_1(1-p_1)+p_2(1-p_2))4\varepsilon]+2\sqrt{2M}\left[(p_1+p_2-2p_1p_2)\sigma(X')\right],
\end{aligned} \tag{29}$$

then, by (25) and (28) we get that (27) holds, and so does (26) as well, as required. Additionally, for large enough M, we neglect all the terms that are not $O(M^2)$ in (29), and thus we require

$$0.87\left(M^2E(X')-2\varepsilon M^2\right) \geq \max_{(p_1,p_2)\in U\backslash U_\delta} M^2[(p_1+p_2-2p_1p_2)(E(X')+2\varepsilon)+(p_1(1-p_1)+p_2(1-p_2))4\varepsilon]. \tag{30}$$

To sum up, we showed in (26) that if δ is such that $0.87S(0,1) > \max_{(p_1,p_2)\in U\backslash U_\delta} S(p_1,p_2)$, then with high probability step 23 of Algorithm 1 returns a δ-precise partition. We also showed that if (30) holds then $0.87S(0,1) > \max_{(p_1,p_2)\in U\backslash U_\delta} S(p_1,p_2)$ holds with high probability. Thus for any δ such that (30) holds, step 23 of Algorithm 1 returns a δ-precise partition with high probability. To find such a δ, we rewrite (30) as

$$0.87E(X') \geq \max_{(p_1,p_2)\in U\backslash U_\delta} (p_1+p_2-2p_1p_2)E(X') + \varepsilon(6(p_1+p_2)+4(p_1^2+p_2^2-p_1p_2)+1.74),$$

and since $(6(p_1+p_2)+4(p_1^2+p_2^2-p_1p_2)+1.74) \leq 21$, we get,

$$0.87E(X') \geq \max_{(p_1,p_2)\in U\setminus U_\delta} (p_1+p_2-2p_1p_2)E(X') + 21\varepsilon,$$

or

$$0.87 - \frac{21\varepsilon}{E(X')} \geq \max_{(p_1,p_2)\in U\setminus U_\delta} (p_1+p_2-2p_1p_2),$$

or using Lemma 9 in Appendix A,

$$0.87 - \frac{63}{4}\varepsilon \geq \max_{(p_1,p_2)\in U\setminus U_\delta} (p_1+p_2-2p_1p_2). \tag{31}$$

That is, for any δ such that (31) holds, for large enough M, step 23 of Algorithm 1 returns with high probability a δ-precise partition.

It is easy to see that the set

$$\mathcal{U}_\delta = [0,1]^2 \setminus \left\{(p_1,p_2) \in [0,1]^2 \mid p_1 \geq \delta, 1-p_2 \geq \delta \text{ or } 1-p_1 \geq \delta, p_2 \geq \delta\right\}$$

satisfies that $U\setminus U_\delta \subset \mathcal{U}_\delta$, that the maximum of $p_1+p_2-2p_1p_2$ on $\mathcal{U}_\delta$ is achieved on the boundary, and that the maximal value is δ. Thus, $\max_{(p_1,p_2)\in U\setminus U_\delta}(p_1+p_2-2p_1p_2) \leq \delta$. Thus, for any δ such that $0.87 - \frac{63}{4}\varepsilon \geq \delta$ we have that (31) holds, and so for large enough N, with high probability, we have a δ-precise partition. The largest δ for which the latter condition holds is $\delta = 0.87 - \frac{63}{4}\varepsilon$, and so step 23 of Algorithm 1 returns a partition which is at least $0.87 - \frac{63}{4}\varepsilon$ precise. □

6 Experimental Results

In this section we show results of Algorithm 1 on both simulated and experimental data sets. For simulated data, it is easy to measure the accuracy of the algorithm for different noise levels. However, in the case of experimental data, since we do not know the correct class partition, the quality of the results is assessed by the Fourier shell correlation [26] of the reconstructed volumes.

All the experiments were executed in MATLAB on a computer with two Intel Xeon X5560 CPUs running at 2.8GHz and an nVidia GTX TITAN GPU. The total running time for the whole algorithm running on 5000 images was around 50 hours.

6.1 Implementation Notes

While in the analysis of Algorithm 1 we use the Goemans-Williamson algorithm for the max-cut problem, for which a performance bound can be proven, this algorithm has memory complexity of $O(N^4)$. Thus, for large N (larger than 1000) it is impractical. Instead, we use a simpler algorithm for the max-cut problem that starts with 8 initial guesses for the cut, makes a local search around

each guess, and returns the best cut detected. Although this naive algorithm does not guarantee an 87% approximation bound as the Goemans-Williamson algorithm, it is no worse in practice.

The running time of Algorithm 1 depends on the running times of the LUD algorithm, the max-cut algorithm, and the number of iterations used. In the experiments presented below, we used 8 iterations of the algorithm, which was more than enough for convergence.

6.2 Simulated data sets with two molecules

To evaluate the performance of Algorithm 1, we first applied it to simulated heterogeneous data sets at various levels of noise. The heterogeneous data sets for the experiment were generated as follows. First, we created two three-dimensional volumes (molecules) corresponding to two types of molecules. The first molecule was a known density map of the 50S subunit of the E. coli ribosome, and the second was its perturbed version created by adding a small sphere. The two three-dimensional volumes were chosen deliberately to have similar structures. Visualizations of the three-dimensional volumes are given in Figure 2. Then, we generated 2500 noiseless projections of each of the volumes, using uniformly distributed random orientations. The resulting 5000 projections, each of size 65×65 pixels, consisted our noiseless heterogeneous data set. Finally, for each level of noise, we added to each image in the noiseless data set additive white Gaussian noise at the given noise level, and applied our algorithm to the resulting noisy data set. As mentioned above, we used 8 iterations in our algorithm. Note that in this case the two classes have equal size with $N_1 = N_2 = 2500$.

The results of the algorithm are summarized in Table 1. Each row in the table corresponds to an experiment at a fixed SNR, and shows the number of images assigned to each class, the parentage of correctly detected common lines (defined as common lines that deviate by up to 10° from the true common lines known from the simulation), and the precision of the partition measured according to Definition 7. To illustrate the SNR values used, we show in Figure 3 a clean image and its noisy realizations at different levels of noise. In Figure 4 we show the reconstructed volumes using the noisy images and the estimated rotations and class partitions.

In the next experiment, we used the same setup as above, except that this time the classes were unbalanced, with $N_1 = 4000$ and $N_2 = 1000$. The results of this experiment are summarized in Table 2. One can see from Tables 1 and 2 that the two classes returned by Algorithm 1 tend to be of similar size. We can see in Table 2 that the estimated class 2 is always larger than its true size, and for higher noise levels it is very clear that the estimated classes tend to be of similar sizes. The intuition for this behavior is that for classes of similar size, the number of edges in the cut is maximal. Moreover, this behavior is advantageous when the number of images of one of the molecules is much larger than of the other (e.g., without loss of generality $N_1 \gg N_2$). In this case, for high noise levels, since the partition will result in two classes of comparable sizes, the estimated class 2 will contain many incorrectly assigned images. However, in class 1 there will be less than

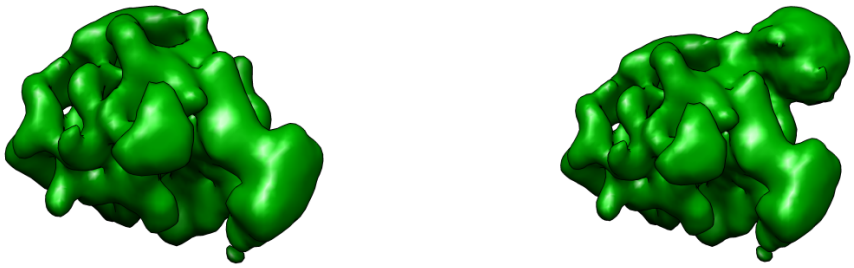

Figure 2: Three-dimensional view of the two volumes considered in the simulated data experiment.

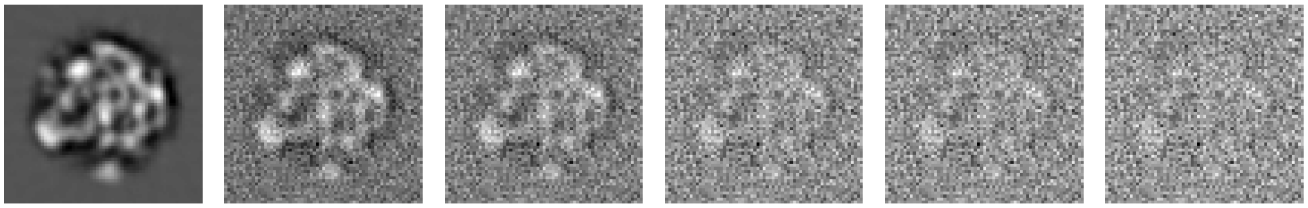

(a) Without noise (b) SNR = 1 (c) SNR = 0.5 (d) SNR = 0.15 (e) SNR = 0.1 (f) SNR = 0.05

Figure 3: A clean projection image and its noisy relations at different levels of noise. Each image is of size 65×65 pixels.

SNR		Correct class		Class size	Precision	% Correct common lines
		class 1	class 2			
1	class 1	2500	0	2500	1	91.48%
	class 2	0	2500	2500		
0.5	class 1	2499	0	2499	0.9996	73.62%
	class 2	1	2500	2501		
0.15	class 1	2411	4	2418	0.9656	38.80%
	class 2	89	2496	2585		
0.1	class 1	2319	163	2482	0.9281	27.77%
	class 2	181	2337	2518		
0.05	class 1	2176	234	2410	0.8749	11.50%
	class 2	324	2266	2590		
0.02	class 1	1604	774	2378	0.6583	4.36%
	class 2	896	1726	2622		
		2500	2500			

Table 1: Results of Algorithm 1 for balanced classes.

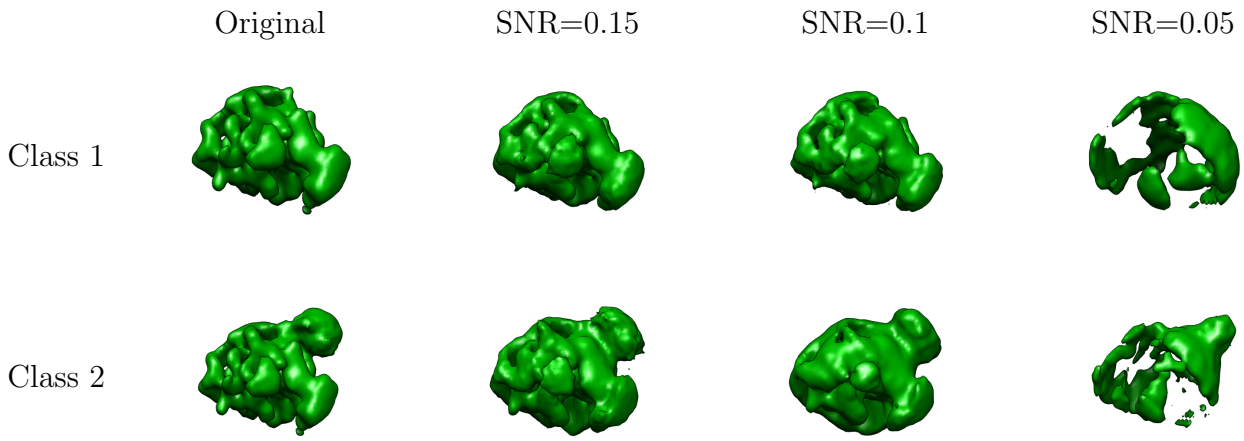

Figure 4: The two volumes reconstructed from the heterogeneous data set at different levels of noise.

N_1 images, yet not many incorrect ones.

6.3 Real data set

For a real data experiment, we used the data [1] provided by the EMDB test image data [2]. This data set consists of 10,000 images of the 70S subunit of the E. coli ribosome with and without Elongation factor G (EF-G). Each image is of size 130×130 pixels, with pixel size of 2.82Å. The data set was preprocessed by the ASPIRE software package [3], by downsampling all images to size 89×89 pixels (in order to improve their SNR), prewhitening the background noise in each image, and normalizing the noise in each image to zero mean and unit variance. After preprocessing, we randomly split the 10,000 images into two disjoint groups of 5000 images each. The reconstructions from these two disjoint groups would be compared below for validating their correctness. Next, for each group independently, we estimated class averages by averaging each image with its 20 nearest images (after proper rotational and translational alignment). A sample of the class averages from each of the two groups in shown in Figure 5. From each of the two groups we then selected the best 3000 class averages. We denote the sets of 3000 class averages corresponding to groups 1 and 2 by S_1 and S_2 , respectively. We then applied Algorithm 1 twice – once to the class averages in S_1 and once to the class averages in S_2. In principle, the original data set is expected to contain two rather homogeneous classes. However, since the data set inevitably contains images of poor quality, some of the resulting class averages may be of poor quality as well. In particular, these class averages may not correspond to any of the underlying structures. We therefore applied Algorithm 1 with $K = 3$ classes, corresponding to the two underlying classes as well as a class of poor-quality averages ("garbage" class averages). We show in Table 3 the number of images in each class in the partition of S_1 and S_2. In order to measure the quality of the results, we used

SNR		Correct class		Class size	Precision	% Correct common lines
		class 1	class 2			
1	class 1	3989	0	3989	1	91.47%
	class 2	11	1000	1011		
0.5	class 1	3531	0	3531	1	69.27%
	class 2	469	1000	1469		
0.15	class 1	2609	57	2667	0.978	31.70%
	class 2	1391	943	2333		
0.1	class 1	2578	63	2641	0.976	25.95%
	class 2	1422	937	2359		
0.05	class 1	2498	81	2579	0.969	13.72%
	class 2	1502	919	2421		
0.02	class 1	1991	428	2419	0.823	5.89%
	class 2	572	2009	2581		
		4000	1000			

Table 2: Results of Algorithm 1 for unbalanced classes.

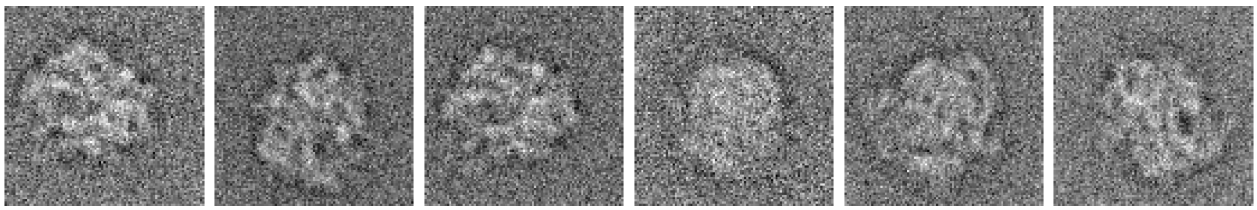

Figure 5: A sample of the class averages.

the Fourier shell correlation (FSC) curves [26] between the reconstructions from the sets S_1 and S_2. As the output of Algorithm 1 for each of the sets S_1 and S_2 consists of three volumes, we show in Figure 6 the FSC curves between the matching pairs of volumes reconstructed from S_1 and S_2 (which are supposed to be two different reconstructions of volumes from the same class), while in Figure 7 we show the FSC curves between volumes corresponding to different classes. In Table 4 we show the FSC between all the reconstructions from S_1 and S_2. As one can see, the matching volumes reconstructed from S_1 and S_2 show greater similarity than the mismatching volumes.

7 Conclusion

We presented a new algorithm for approximating the class partition of a heterogeneous image data set in cryo-EM. We derived theoretical bounds for the algorithm, and applied it on both simulated and experimental data sets. For simulated data, the algorithm finds accurate partitions,

	Number of images		
	class 1	class 2	"garbage" class
S_1	819	953	1228
S_2	844	960	1196

Table 3: Number of images in each class in the results of Algorithm 1 on S_1 and S_2.

	class 1	class 2	"garbage" class
class 1	29.33Å	42.09Å	36.55Å
class 2	34.91Å	27.49Å	33.25Å
"garbage" class	33.92Å	28.63Å	40.61Å

Table 4: FSC between the volumes reconstructed from the experimental data.

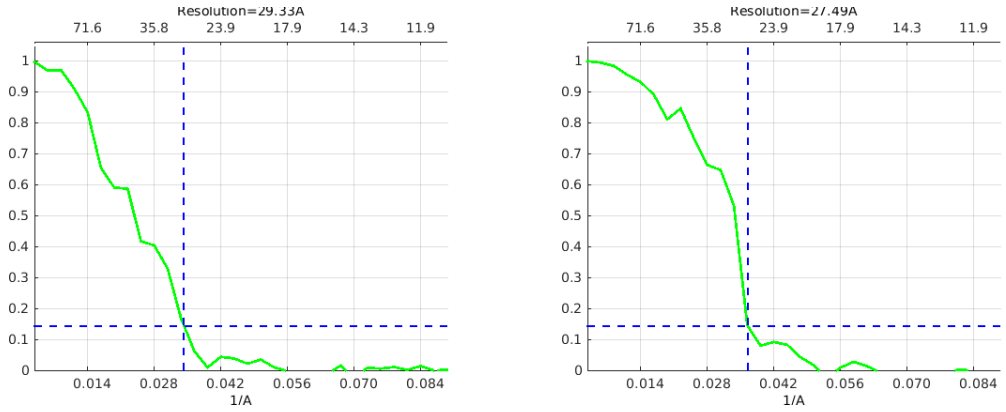

(a) (b)

Figure 6: FSC of matching classes.

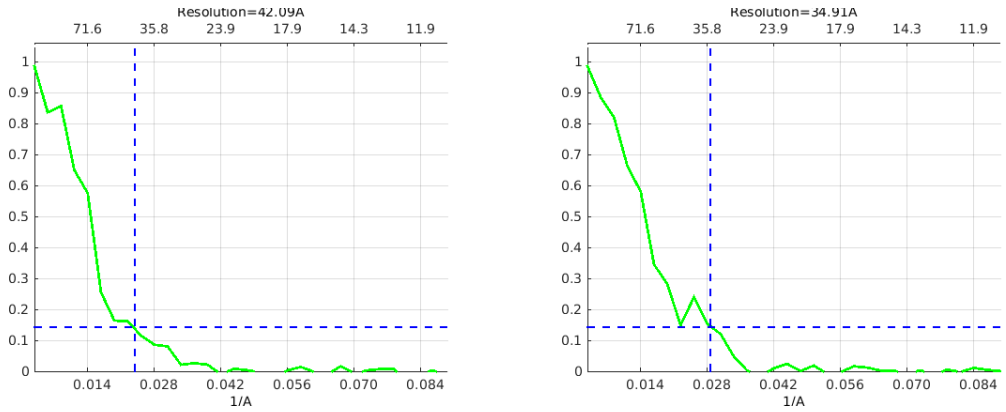

(a) (b)

Figure 7: FSC of mismatching classes.

even under extreme levels of noise. As the proposed algorithm is based on the LUD [27], it is applicable only to molecules without symmetries. However, it can be easily combined with abinitio reconstruction algorithms for molecules with symmetries. Moreover, the LUD can be completely replaced with another orientation assignment algorithm such as [22], in which case the score used for comparing images should be revised. The proposed algorithm can also be incorporated with confidence information regarding the score between each pair of images. This score is available as a byproduct of the algorithms [18, 22], and may further improve the robustness of our algorithm to noise.

Appendix A Distributions

Lemma 9. *Let X be the random variable corresponding to the distance between two uniformly distributed random points on the unit sphere. Then the PDF of X is $f(r) = r/2$, and its expectancy is $E(X) = \frac{4}{3}$.*

The proof is given in [24].

Lemma 10. *Let $Y_i \sim N(\mu_i, \sigma_i)$ be independent normally distributed random variables. Then, for N sufficiently large, $\max_{1\le i\le N} Y_i$ is bounded with high probability by $\max_{1\le i\le N} \mu_i + 2\sqrt{\log N} \max_{1\le i\le N} \sigma_i$.*

Proof. Since $Y_i \sim N(\mu_i, \sigma_i)$ we have that

$$\begin{aligned} P(Y_i > \mu_i + t\sigma_i) &= \frac{1}{\sigma_i\sqrt{2\pi}} \int_{\mu_i+t\sigma_i}^{\infty} e^{-\frac{(x-\mu_i)^2}{2\sigma_i^2}} dx = \frac{1}{\sigma_i\sqrt{2\pi}} \int_{t\sigma_i}^{\infty} e^{-\frac{x^2}{2\sigma_i^2}} dx \\ &= \frac{1}{\sqrt{2\pi}} \int_t^{\infty} e^{-\frac{x^2}{2}} dx \le \frac{1}{\sqrt{2\pi}} \int_t^{\infty} \frac{x}{t} e^{-\frac{x^2}{2}} dx = \frac{1}{\sqrt{2\pi}} \frac{1}{t} e^{-\frac{t^2}{2}}. \end{aligned}$$

Thus,

$$\begin{aligned} P\left(Y_1 < \mu_1 + t\sigma_1, \ldots, Y_N < \mu_N + t\sigma_N\right) &= \prod_{i=1}^{N} P(Y_i < \mu_i + t\sigma_i) = \left(1 - \frac{1}{\sqrt{2\pi}} \frac{1}{t} e^{-\frac{t^2}{2}}\right)^N \\ &> 1 - N\frac{1}{\sqrt{2\pi}} \frac{1}{t} e^{-\frac{t^2}{2}} = 1 - \frac{1}{\sqrt{2\pi}} \frac{1}{t} e^{n\log_2 e - \frac{t^2}{2}}, \end{aligned} \tag{32}$$

where the inequality in (32) follow from Bernoulli's inequality. In particular, for $t = 2\sqrt{\log N}$ we have

$$P(\max_{1\le i\le N} Y_i < \max_{1\le i\le N} \mu_i + 2\sqrt{\log n}\sigma_i) \ge 1 - \frac{1}{\sqrt{2\pi}} \frac{1}{2\sqrt{\log N}} e^{\log N(\log_2 e - 2)}.$$

Thus,

$$\lim_{N\to\infty} P(\max_{1\le i\le N} Y_i > \max_{1\le i\le N} \mu_i + 2\sqrt{\log N}\sigma_i) \to 0.$$

□

References

[1] `ftp://ftp.ebi.ac.uk/pub/databases/emtest/SPIDER_FRANK_data/`.

[2] `https://www.ebi.ac.uk/pdbe/emdb/test_data.html`.

[3] Algorithms for single particle reconstruction. `http://spr.math.princeton.edu/`.

[4] Bai, X. and McMullan, G. ans Scheres, S. H. W. How cryo-EM is revolutionizing structural biology. *Trends in Biochemical Sciences*, 40(1):49—-57, 2015.

[5] Bartesaghi, A., Merk, A., Banerjee, S., Matthies, D., Wu, X., Milne, J. L. S., and Subramaniam, S. 2.2Å resolution cryo-EM structure of β-galactosidase in complex with a cell-permeant inhibitor. *Science*, 348(6239):1147–1151, 2015.

[6] Callaway, E. The revolution will not be crystallized: a new method sweeps through structural biology. *NATURE*, 525(7568):172–174, 2015.

[7] Cheng, Y., Grigorieff, N., Penczek, P. A., and Walz, T. A primer to single-particle cryo-electron microscopy. *Cell*, 161(3):438—-449, 2015.

[8] Frank, J. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. Oxford, 2006.

[9] Frieze, A. and Jerrum, M. Improved approximation algorithms for MAX k-CUT and MAX BISECTION. In *Integer Programming and Combinatorial Optimization*, pages 1–13. Springer, 1995.

[10] Goemans, M. X. and Williamson, D. P. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

[11] Herman, G. T. *Fundamentals of Computerized Tomography: Image Reconstruction from Projections*. Springer, London, UK, 2nd edition, 2009.

[12] Herman, G. T. and Kalinowski, M. Classification of heterogeneous electron microscopic projections into homogeneous subsets. *Ultramicroscopy*, 108(4):327–338, 2008.

[13] Katsevich, G., Katsevich, A., and Singer, A. Covariance matrix estimation for the cryo-em heterogeneity problem. *arXiv preprint arXiv:1309.1737*, 2013.

[14] Liao, H. Y., Hashem, Y., and Frank, J. Efficient estimation of three-dimensional covariance and its application in the analysis of heterogeneous samples in cryo-electron microscopy. *Structure*, 23(6):1129—1137, 2015.

[15] Natterer, F. *The Mathematics of Computerized Tomography*. Classics in Applied Mathematics. SIAM, 2001.

[16] Natterer, F. and Wübbeling, F. *Mathematical Methods in Image Reconstruction*. SIAM Monographs on Mathematical Modeling and Computation. SIAM, 2001.

[17] Penczek, P. A, Frank, J., and Spahn, C. M. T. A method of focused classification, based on the bootstrap 3D variance analysis, and its application to EF-G-dependent translocation. *Journal of structural biology*, 154(2):184–194, 2006.

[18] Pragier, G., Greenberg, I., Cheng, X., and Shkolnisky, Y. A graph partitioning approach to simultaneous angular reconstitution. *IEEE Transactions Computational Imaging*, Accepted for publication.

[19] Scheres, S. H. W. Maximum-likelihood methods in cryo-EM. Part II: Application to experimental data. *Methods in enzymology*, 482:295, 2010.

[20] Scheres, S. H. W. Relion: Implementation of a bayesian approach to cryo-EM structure determination. *Journal of Structural Biology*, 180(3):519—-530, 2012.

[21] Shatsky, M., Hall, R. J., Nogales, E., Malik, J., and Brenner, S. E. Automated multi-model reconstruction from single-particle electron microscopy data. *Journal of Structural Biology*, 170(1):98–108, 2010.

[22] Shkolnisky, Y. and Singer, A. Viewing directions estimation in cryo-EM using synchronization. *SIAM Journal on Imaging Sciences*, 5(3):1088–1110, 2012.

[23] Singer, A. and Shkolnisky, Y. Three-dimensional structure determination from common lines in cryo-EM by eigenvectors and semidefinite programming. *SIAM Journal on Imaging Sciences*, 4(2):543–572, 2011.

[24] Solomon, H. *Geometric probability*, volume 28 of *CBMS-NSF Regional Conference Series in Applied Mathematics*, chapter 6, page 162. SIAM, 1978.

[25] van Heel, M. and Frank, J. Use of multivariate statistics in analysing the images of biological macromolecules. *Ultramicroscopy*, 6(2):187–194, 1981.

[26] van Heel, M. and Schatz, M. Fourier shell correlation threshold criteria. *J. Struct. Biol.*, 151(3):250–262, 2005.

[27] Wang, L., Singer, A., and Wen, Z. Orientation determination of Cryo-EM images using least unsquared deviations. *SIAM journal on imaging sciences*, 6(4):2450–2483, 2013.